Negotiating Socially Optimal Allocations of Resources

Ulle Endriss
Imperial College London

Acknowledgements: Joint work with Nicolas Maudet, Fariba Sadri, Francesca Toni, Yann Chevaleyre and Sylvia Estivie. Funded by the European Commission as part of the SOCS project (IST-2001-32530).

Introduction

- Multiagent systems may be thought of as *"societies of agents"*.
- Agents *negotiate* deals to exchange resources to benefit either themselves or society as a whole.
- Agents may use very simple rationality criteria to decide what deals to accept, but interaction patterns may be complex (*multilateral* deals).

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Resource Allocation by Negotiation

- Finite set of *agents* $\mathcal{A}$ and finite set of discrete *resources* $\mathcal{R}$.
- An *allocation* A is a partitioning of $\mathcal{R}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_3, r_7\}$ — agent i owns resources r_3 and r_7
- Every agent $i \in \mathcal{A}$ has got a *utility function* $u_i : 2^{\mathcal{R}} \to \mathbb{Q}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy
- Agents may engage in negotiation to exchange resources in order to benefit either themselves or society as a whole.
- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may come with a number of side payments to compensate some of the agents for a loss in utility. A *payment function* is a function $p : \mathcal{A} \to \mathbb{Q}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.
  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* £5, while agent j *receives* £5.

Individual Rationality

A *rational* agent (who does not plan ahead) will only accept deals that improve its individual welfare:

> *A deal* $\delta = (A, A')$ *is called individually rational iff there exists a payment function* p *such that* $u_i(A') - u_i(A) > p(i)$ *for all* $i \in \mathcal{A}$*, except possibly* $p(i) = 0$ *for agents* i *with* $A(i) = A'(i)$*.*

That is, an agent will only accept a deal *iff* it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility).

Utilitarian Social Welfare

The *social welfare* associated with an allocation of resources A is defined as follows:

$$sw(A) \quad = \quad \sum_{i \in \mathcal{A}gents} u_i(A)$$

This is the so-called *utilitarian* definition of social welfare, which measures the "sum of all pleasures" (Jeremy Bentham, ∼1820).

▶ Observe that maximising this function amounts to maximising the *average utility* enjoyed by agents in the system.

Example

Let $\mathcal{A} = \{ann, bob\}$ and $\mathcal{R} = \{chair, table\}$ and suppose our agents use the following utility functions:

$$\begin{array}{rclrcl}
u_{ann}(\{\,\}) & = & 0 & u_{bob}(\{\,\}) & = & 0 \\
u_{ann}(\{chair\}) & = & 2 & u_{bob}(\{chair\}) & = & 3 \\
u_{ann}(\{table\}) & = & 3 & u_{bob}(\{table\}) & = & 3 \\
u_{ann}(\{chair, table\}) & = & 7 & u_{bob}(\{chair, table\}) & = & 8
\end{array}$$

Furthermore, suppose the initial allocation of resources is A_0 with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \{\,\}$.

▶ Social welfare for allocation A_0 is 7, but it could be 8. By moving only a *single* resource from agent ann to agent bob, the former would lose more than the latter would gain (not individually rational). The only possible deal would be to move the whole *set* $\{chair, table\}$.

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Linking the Local and the Global Perspectives

Lemma 1 (Individual rationality) *A deal $\delta = (A, A')$ with side payments is individually rational iff $sw(A) < sw(A')$.*

Proof. "$\Rightarrow$": Use definitions.
"$\Leftarrow$": Every agent will get a positive payoff if the following payment function is used:

$$p(i) \;=\; u_i(A') \;-\; u_i(A) \;-\; \underbrace{\frac{sw(A') - sw(A)}{|\mathcal{A}|}}_{>0} \qquad \Box$$

▶ We can now prove a first result on negotiation processes:

Lemma 2 (Termination) *There can be no infinite sequence of individually rational deals, i.e. negotiation must always terminate.*

Proof. The space of distinct allocations is finite and, by Lemma 1, every rational deal results in a strict increase in social welfare. $\Box$

Negotiating Socially Optimal Allocations

The following result is due to Sandholm (originally for distributed task allocation problems):

Theorem 3 (Convergence) *Any sequence of individually rational deals will eventually result in an allocation with maximal social welfare.*

Proof. By Lemma 2, negotiation must terminate. Assume the final allocation A is *not* optimal, i.e. there exists an allocation A' with $sw(A) < sw(A')$. But then, by Lemma 1, the deal $\delta = (A, A')$ would be individually rational (contradicts assumption of A being final). $\square$

▶ Agents can act *locally* and need not be aware of the global picture (convergence towards a global optimum is guaranteed by the theorem).

T. Sandholm. *Contract types for satisficing task allocation: I Theoretical results.* AAAI Spring Symposium 1998.

Multilateral Negotiation

Optimal outcomes can only be guaranteed if the negotiation protocol allows for deals involving *any number of agents* and *resources:*

Theorem 4 (Necessity of complex deals) *Any deal $\delta = (A, A')$ may be necessary, i.e. there are utility functions and an initial allocation such that any sequence of individually rational deals leading to an allocation with maximal social welfare would have to include δ.*

Proof. By systematic definition of utility functions such that A' is optimal and A is second best ... □

▶ Note that most work on negotiation in multiagent systems is on *bilateral* ("one-to-one") negotiation ... or on *auctions*.

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Negotiation without Money

- Problem: Agents may require *unlimited amounts of money* to get through a negotiation process.
- Without side payments, however, rational negotiation cannot guarantee outcomes with maximal social welfare.
  Example: *Would you give me your car just because I value it more highly than you do? ... note that this would be socially beneficial!*
- We have results that show that *cooperatively rational* deals (only one agent requires a *strictly* positive payoff) without side payments are sufficient to negotiate *Pareto optimal* allocations (and multilateral deals are again necessary).
- Future work: Study in more detail what happens in scenarios with *limited amounts* of money of *limited granularity*.

U. Endriss, N. Maudet, F. Sadri and F. Toni. *On optimal outcomes of negotiations over resources*. AAMAS-2003.

Restricted Domains

It is difficult to design protocols for truly multilateral negotiation, but in restricted domains simple protocols can sometimes suffice:

Theorem 5 (Additive domains) *If all utility functions are additive, then individually rational one-resource deals with side payments suffice to guarantee outcomes with maximal social welfare.*

Proof. Use the fact that $\sum_{\mathcal{A}}\sum_{\mathcal{R}}$ can be rewritten as $\sum_{\mathcal{R}}\sum_{\mathcal{A}}$ to find an applicable one-resource deal for any sub-optimal allocation ... $\square$

- 0-1 domains: All agents use additive utility functions assigning 1 (want it) or 0 (don't want it) to single resources.

Theorem 6 (0-1 domains) *In 0-1 domains, even cooperatively rational one-resource deals without side payments suffice.*

Proof. Similar. $\square$

Alternative Representation of Utility Functions

- Problem: The *"bundle form"* of representing utility functions can be problematic if there are too many bundles with non-zero values.

- A utility function is called *k-additive* *iff* the utility assigned to a bundle R can be represented as the sum of basic utilities assigned to subsets of R with cardinality $\leq k$ (*limited synergies*).

- The *k-additive form* of representing utility functions:

$$u_i(R) \quad = \sum_{T \subseteq \mathcal{R},\, |T| \leq k} \alpha_i^T \times I_R(T) \;\text{ with } I_R(T) = \begin{cases} 1 & \text{if } T \subseteq R \\ 0 & \text{otherwise} \end{cases}$$

  Example: $u_i = 3.r_1 + 7.r_2 - 2.r_2.r_3$ is a 2-additive function

- Note that *any* utility function is representable as a k-additive function for some $k \leq |\mathcal{R}|$.

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Multiagent resource allocation with k-additive utility functions*. DIMACS-LAMSADE Workshop 2004.

Separation Results

Theorem 7 (Efficiency of the k-additive form) *The bundle form cannot polynomially simulate the k-additive form.*

Proof. Consider the utility function $u_i(R) = |R|$. Representing u_i requires $|\mathcal{R}|$ non-zero coefficients in the k-additive form (*linear*), but $2^{|\mathcal{R}|}-1$ non-zero values in the bundle form (*exponential*). $\square$

Theorem 8 (Efficiency of the bundle form) *The k-additive form cannot polynomially simulate the bundle form.*

Proof. Consider the utility function $u_i(R) = \begin{cases} 1 & \text{if } |R| = 1 \\ 0 & \text{otherwise} \end{cases}$

Requires $|\mathcal{R}|$ non-zero values in the bundle form (*linear*), but $2^{|\mathcal{R}|}-1$ non-zero coefficients in the k-additive form (*exponential*): namely $\alpha_i^T = 1$ for $|T| = 1$, $\alpha_i^T = -2$ for $|T| = 2$, $\alpha_i^T = 3$ for $|T| = 3$, ... $\square$

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Connections to Combinatorial Auctions

While our negotiation framework is clearly *not an auction*, there are still interesting connections between combinatorial auctions and the "centralised problem" of finding a socially optimal allocation:

Combinatorial auctions	**Negotiation**
Bidders submitting (several) bids	agents with utility functions
Bidding language (XOR)	representation of utilities (bundle form)
Revenue for the auctioneer	utilitarian social welfare
Winner determination problem	finding an optimal allocation
(Usually) free disposal	no free disposal (depends on agents)

▶ The "standard" bidding language (OR language) is less expressive than either the bundle or the k-additive form.

▶ I'm not aware of work on combinatorial auctions using a bidding language corresponding to the k-additive form (could be interesting).

Complexity of Maximising Social Welfare

Winner determination (more precisely: the underlying decision problem) in combinatorial auctions is known to be NP-complete. So this is not a surprising result:

Theorem 9 (Complexity) *The decision problem underlying the problem of finding an allocation with maximal utilitarian social welfare is NP-complete (wrt. the representation of utilities in bundle form).*

Proof. (i) NP-membership: for any proposed allocation A, we can check $sw(A) > K$ in polynomial time ✓

(ii) NP-hardness: by reduction from $\textsc{Weighted Set Packing}$ (see next slide for details) $\square$

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Multiagent resource allocation with k-additive utility functions*. DIMACS-LAMSADE Workshop 2004.

Proof of NP-hardness

The following problem is known to be NP-complete:

WEIGHTED SET PACKING
Instance: Collection $\mathcal{C}$ of finite sets with positive weights.
Solution: Collection of disjoint sets $\mathcal{C}' \subseteq \mathcal{C}$.
Question: Does the sum of weights of the sets in $\mathcal{C}'$ exceed K?

This can be reduced to our problem as follows:

- For every set R in $\mathcal{C}$ with weight x, introduce an agent i and define $u_i(R) = x$ and $u_i(R') = 0$ for all bundles $R' \neq R$.
- “Free disposal”: introduce an additional agent i^* with $u_{i^*} \equiv 0$.

Now any allocation A with $sw(A) > K$ corresponds to a set packing $\mathcal{C}'$ with a sum of weights exceeding K. Hence, our problem is at least as hard as WEIGHTED SET PACKING. $\square$

Related Complexity Results

- Dunne *et al.* (2003) show NP-completeness of the same problem, but with respect to different parameters:
  - NP-hardness with respect to the *number of resources*
  - NP-membership with respect to a *compact* representation of utility functions as programs
- We also have an NP-completeness result with respect to utility functions given in k-*additive form*.

P. E. Dunne, M. Wooldridge and M. Laurence. *The complexity of contract negotiation*. Technical Report ULCS-03-002, University of Liverpool 2003.

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Multiagent resource allocation with k-additive utility functions*. DIMACS-LAMSADE Workshop 2004.

Communication Complexity

- Our complexity results so far concern the *computational* complexity of an *abstract* problem: finding a socially optimal allocation *somehow* (not necessarily by negotiation).
- What we are really interested in is the complexity of negotiation processes in our multilateral trading framework.
- We therefore consider also the *communication complexity* of negotiating socially optimal allocations of resources, i.e. we focus on the length of negotiation processes and the amount of information exchanged, rather than on computational aspects.

U. Endriss and N. Maudet. *On the communication complexity of multilateral trading*. AAMAS-2004.

Aspects of Complexity

(1) How many *deals* are required to reach an optimal allocation?

- communication complexity as number of individual deals
- technical results to follow

(2) How many *dialogue moves* are required to agree on one such deal?

- affects communication complexity as number of dialogue moves

(3) How expressive a *communication language* do we require?

- Minimum requirements: performatives *propose*, *accept*, *reject* + content language to specify multilateral deals
- affects communication complexity as number of bits exchanged

(4) How complex is the *reasoning* task faced by an agent when deciding on its next dialogue move?

- computational complexity (local rather than global view)

Number of Deals

We have two results on *upper bounds* pertaining to the first variant of our negotiation framework (with side payments, general utility functions, and aiming at maximising utilitarian social welfare):

Theorem 10 (Shortest path) *A single rational deal is sufficient to reach an allocation with maximal social welfare.*

Proof. Use Lemma 1 [$\delta = (A, A')$ rational iff $sw(A) < sw(A')$]. $\square$

Theorem 11 (Longest path) *A sequence of rational deals can consist of up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals, but not more.*

Proof. No allocation can be visited twice (same lemma) and there are $|\mathcal{A}|^{|\mathcal{R}|}$ distinct allocations $\Rightarrow$ upper bound follows $\checkmark$
To show that the upper bound is *tight*, we need to show that it is possible that all allocations have distinct social welfare (see paper). $\square$

Further Results

- Number of rational deals without side payments required to reach a Pareto optimal allocation of resources:
  - *Shortest path:* ≤ 1
  - *Longest path:* $< |\mathcal{A}| \cdot (2^{|\mathcal{R}|} - 1)$
- Number of rational one-resource deals with side payments to reach an allocation with maximal social welfare in additive domains:
  - *Shortest path:* $\leq |\mathcal{R}|$
  - *Longest path:* $\leq |\mathcal{R}| \cdot (|\mathcal{A}| - 1)$
- Number of rational one-resource deals without side payments to reach an allocation with maximal social welfare in 0-1 domains:
  - *Shortest and longest path:* $\leq |\mathcal{R}|$

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Egalitarian Agent Societies

- The utilitarian sw is not the only *collective utility function* ...
- The *egalitarian* collective utility function sw_e, for instance, measures social welfare as follows:

$$sw_e(A) \quad = \quad min\{u_i(A) \mid i \in \mathcal{A}gents\}$$

Maximising this function amounts to improving the situation of the weakest members of society.

- We have defined a local rationality criterion ("equitable deals") for agents operating in egalitarian systems and proved convergence and necessity theorems similar to those we have seen earlier.

U. Endriss, N. Maudet, F. Sadri and F. Toni. *Resource allocation in egalitarian agent societies*. MFI-2003.

Utilitarianism versus Egalitarianism

- In the MAS literature the utilitarian viewpoint (that is, social welfare = sum of individual utilities) is usually taken for granted.
- In philosophy/sociology/economics not.
- John Rawls' *"veil of ignorance"* (*A Theory of Justice*, 1971):

> *Without knowing what your position in society (class, race, sex, ...) will be, what kind of society would you choose to live in?*

- Reformulating the *veil of ignorance for multiagent systems:*

> *If you were to send a software agent into an artificial society to negotiate on your behalf, what would you consider acceptable principles for that society to operate by?*

- Conclusion: worthwhile to investigate egalitarian (and other) social principles also in the context of multiagent systems.

Notions of Social Welfare

- *Utilitarian:* sum of utilities $sw_u(A) = \sum_{i \in \mathcal{A}} u_i(A)$
- *Nash product:* product of utilities $sw_N(A) = \prod_{i \in \mathcal{A}} u_i(A)$
- *Egalitarian:* utility of the weakest $sw_e(A) = min\{u_i(A) \mid i \in \mathcal{A}\}$
- *Elitist:* utility of the strongest $sw_{el}(A) = max\{u_i(A) \mid i \in \mathcal{A}\}$
- *Pareto optimality:* no other allocation is better for some agents without being worse for others
- *Lorenz optimality:* the sum of utilities of the k weakest agents cannot be maintained for all and increased for some $k \leq |\mathcal{A}|$
- *Envy-freeness:* no agent would rather have the bundle allocated to one of the other agents $u_i(A(i)) \geq u_i(A(j))$
  - envy-free allocations are not always possible
  - could search for *envy-reducing* deals (for instance, with respect to the number of envious agents or the average degree of envy)

Welfare Engineering

- Choice (and possibly design) of *social welfare orderings* that are appropriate for specific agent-based applications.
  - Example: The *elitist* collective utility function sw_{el} seems unethical for human society, but may be appropriate for a distributed application where each agent gets the same task.
  - Slogan: "welfare economics for *artificial* agent societies"
- Design of suitable *rationality criteria* for agents participating in negotiation in view of different notions of social welfare.
  - Example: To achieve *Lorenz optimal* allocations in *0-1 domains without money*, ask agents to negotiate *cooperatively rational* or *inequality-reducing* deals over *one resource* at a time.
  - Slogan: "*inverse* welfare economics" ($\leadsto$ mechanism design)

U. Endriss and N. Maudet. *Welfare engineering in multiagent systems*. ESAW-2003.

Criteria for Social Welfare Choice

We have tried to identify criteria that determine what social welfare ordering is appropriate for which application (work in progress):

- What does the income of the system provider depend on?
  - *Utility-dependent* ("tax on gain") $\leadsto$ utilitarian
  - *Membership-dependent* ("joining fee") $\leadsto$ "fair" approach
  - *Transaction-dependent* ("pay as you go") $\leadsto$ not clear (but note the connections to communication complexity)
- Can agents join or leave the society *during* negotiation? Yes: review definitions (e.g. utilitarian welfare as average utility)
- Can agents participate in *more than one* negotiation? Yes: strong point for fair approaches (egalitarian, envy-reducing)

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Welfare engineering in practice: On the variety of multiagent resource allocation problems*. ESAW-2004.

Talk Outline

- Resource allocation by negotiation in multiagent systems
  - *definition of our basic negotiation framework*
- Fundamental results for the basic framework
  - *links between individual rationality and social welfare*
  - *convergence to optimal states and the need for multilateral deals*
- Variations on the basic framework
  - *the problem of unlimited money and results without money*
  - *restricted domains and alternative representations of utility functions*
- Complexity issues
  - *connections to combinatorial auctions*
  - *computational complexity and communication complexity*
- Welfare engineering
  - *the veil of ignorance in multiagent systems*
  - *egalitarian and elitist agent societies, envy-free allocations, . . .*
- Conclusions

Conclusions

- Multilateral negotiation is an exciting and fruitful area of research:
  - Knowledge transfer from economics to computer science and AI
  - Application of computational tools to problems in economics
- Many open problems and scope for new directions of research:
  - *Conceptual:* "ethics of multiagent systems"?
  - *Methodological:* turn the basic ideas of welfare engineering into a practical design methodology for agent-based systems
  - *Practical:* design multilateral negotiation protocols
  - *Technical:* complexity issues and the like
  - *Algorithmic:* use optimisation algorithms to guide negotiation
  - *Experimental:* simulate negotiation and develop heuristics
- Papers are available from my website:

`http://www.doc.ic.ac.uk/~ue/`